

Advanced Electronic Structure Methods for Heterogeneous Catalysis and Separation of Heavy Metals

Technical Report for the ALCF Theta Early Science Program

Argonne Leadership Computing Facility

ALCF Early Science Program (ESP) Technical Report

ESP Technical Reports describe the code development, porting, and optimization done in preparing an ESP project's application code(s) for the next generation ALCF computer system. This report is for a project in the Theta ESP, preparing for the ALCF Theta computer system.

About Argonne National Laboratory

Argonne is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC under contract DE-AC02-06CH11357. The Laboratory's main facility is outside Chicago, at 9700 South Cass Avenue, Argonne, Illinois 60439. For information about Argonne and its pioneering science and technology programs, see www.anl.gov.

DOCUMENT AVAILABILITY

Online Access: U.S. Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available free at OSTI.GOV (<http://www.osti.gov/>), a service of the U.S. Dept. of Energy's Office of Scientific and Technical Information

Reports not in digital format may be purchased by the public from the National Technical Information Service (NTIS):

U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Rd
Alexandria, VA 22312
www.ntis.gov
Phone: (800) 553-NTIS (6847) or (703) 605-6000
Fax: (703) 605-6900
Email: orders@ntis.gov

Reports not in digital format are available to DOE and DOE contractors from the Office of Scientific and Technical Information (OSTI):

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
www.osti.gov
Phone: (865) 576-8401
Fax: (865) 576-5728
Email: reports@osti.gov

Disclaimer

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor UChicago Argonne, LLC, nor any of their employees or officers, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of document authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, Argonne National Laboratory, or UChicago Argonne, LLC.

Advanced Electronic Structure Methods for Heterogeneous Catalysis and Separation of Heavy Metals

Technical Report for the ALCF Theta Early Science Program

edited by
Timothy J. Williams and Ramesh Balakrishnan

Argonne Leadership Computing Facility

prepared by
Mark S. Gordon, Spencer R. Pruitt, Sarom S. Leang, Graham Fletcher, and Yuri Alexeev

September 2017

Advanced Electronic Structure Methods for Heterogeneous Catalysis and Separation of Heavy Metals

Mark S. Gordon^{*1}, Spencer R. Pruitt^{†2}, Sarom S. Leang^{‡1}, Graham Fletcher^{§3}, and Yuri Alexeev^{¶3}

¹Department of Chemistry, Iowa State University, Ames, IA

²Academic and Research Computing, Worcester Polytechnic Institute, Worcester, MA

³Leadership Computing Facility, Argonne National Laboratory, Argonne, IL

1 Science Summary

For code optimizations on Theta, graphene sheet systems were used for benchmarking purposes. Graphene sheets are of interest to researchers in the area of microlubricants. The proposed scientific applications for this project were not achievable with the support awarded for this project. The performance improvements on Theta reported in this work will directly support the scientific applications described in our recently submitted 2018 INCITE proposal.

2 Codes, Methods and Algorithms

The Hartree-Fock (HF) method is used to iteratively solve the electronic Schrödinger equation for a many-body system. The resulting electronic energy and electronic wave function are used to compute equilibrium geometries and a variety of molecular properties. The HF solution is a starting point for a wide range of higher accuracy methods. The HF equations are solved numerically by self-consistent Field (SCF) iterations. In the SCF iterations, a Fock matrix which approximates the single-electron energy operator for the system is constructed from a core Hamiltonian matrix, density matrix, and set of electron repulsion integrals (ERIs). The Fock matrix is then diagonalized to provide a set of molecular orbital coefficients. These coefficients are used to compute a new density matrix. A new Fock matrix is then constructed as a function of the new density matrix. Diagonalization of the updated Fock matrix provides a new set of MO coefficients which are used to update the density matrix. This iterative process continues until the root-mean-squared difference of consecutive densities are below a chosen convergence threshold. The core SCF kernel is a four-fold loop over indices i, j, k, l which correspond to groups of basis functions called shells. ERIs are

^{*}mark@si.msg.chem.iastate.edu

[†]srpruitt@wpi.edu

[‡]sarom@si.msg.chem.iastate.edu

[§]fletcher@anl.gov

[¶]yuri@anl.gov

computed at the granularity of shell quartets corresponding to i, j, k, l .

The original MPI-based GAMESS HF implementation exhibits a very high memory footprint, because a number of data structures are replicated across MPI ranks. For architectures with many cores such as KNL, the memory footprint of these replicated data structures severely limits the size of chemical systems which can fit in memory. In this work, the HF implementation in the GAMESS electronic structure package was optimized for the Intel Xeon Phi Knight’s Landing architecture (KNL), targeting the Theta supercomputer system. Two new hybrid MPI/OpenMP HF algorithms were implemented in GAMESS, which reduce the HF memory footprint by as much as 200 times compared to the legacy code. The resulting MPI/OpenMP algorithms were shown to run up to six times faster than the MPI-based algorithm across benchmarks utilizing as many as 192,000 cores on Theta.

3 Code Development

3.1 Hybrid MPI/OpenMP SCF Algorithm - Private Fock Matrix

For the first hybrid algorithm, the Fock matrix is replicated among all ranks and threads (as in the MPI-only implementation). All other large data structures are shared, reducing the memory footprint on the latest Xeon and KNL processors by approximately 5 times. For the core SCF kernel, the master thread of each MPI ranks updates the i index. OpenMP parallelization is used over combined j and k shell loops. By joining these loops, each MPI ranks is responsible for a larger pool of tasks compared to the legacy MPI-only implementation. Just like the MPI-only implementation, the final step is reduction of the Fock matrix over MPI processes.

3.2 Hybrid MPI/OpenMP SCF Algorithm - Shared Fock Matrix

For the second hybrid algorithm, the Fock matrix is shared among threads. Each ERI contributes to six Fock matrix elements at one time, and removing Fock matrix write dependencies is not straightforward. The ERI contribution is added to three column-blocks of the Fock matrix, corresponding to the i, j , and k indices. Each block corresponds to all basis functions for a single shell. In the shared Fock algorithm, a thread-private storage buffer is used to store partial Fock matrix contributions. These contributions are flushed to the full matrices when the shell indices change. The buffers are arranged as two-dimensional arrays, with each column corresponding to a thread and each row corresponding to a basis function. Padding bytes are added to the leading dimension of the ray and chunking is used on the reduction step to prevent false sharing. The access pattern of the Fock matrix by the k index spans only a single Fock matrix element, so the corresponding Fock matrix element can be updated directly. Therefore, thread-private storage is required for only the i and j column-blocks of the Fock matrix. The four-fold SCF loops are arranged into two loops over combined ij and kl indices. MPI parallelization is implemented over the ij loop, and OpenMP parallelization is used for the kl loop. At the end of each kl loop, the i buffer is only flushed if the i buffer is changed. This reduces the number of explicit thread synchronizations required.

3.3 Performance Improvements on Theta

Graphene sheets of dimension 0.5nm, 1.0nm, 1.5nm, 2.0nm, and 5.0nm were used for benchmarking. The 6-31G(d) basis set was used in all calculations. The density, Fock, atomic orbital overlap, one-electron Fock matrices, and molecular orbital coefficients are NxN in size for N basis functions where N ranged from 660 to 30,240 elements for this dataset. The asymptotic memory footprint computed as the number of basis functions approaches infinity is listed for each algorithm in Table 1.

Table 1: Size characteristics and memory requirements for chemical systems used in benchmarks

System	# Basis Func.	MPI Memory, GB	Private Fock Memory, GB	Shared Fock Memory, GB
0.5nm	660	7	0.13	0.03
1.0nm	1800	48	1	0.2
1.5nm	3300	160	3	0.8
2.0nm	5340	417	8	2
5.0nm	30240	9869	257	52

The memory and cluster modes of the KNL processor were benchmarked for the three algorithms using the 0.5nm and 2.0nm datasets (Figure 1). The smaller dataset indicates a higher sensitivity to the memory and cluster modes compared to the larger dataset. With the exception of All-to-All mode with the 0.5nm dataset, the shared Fock version outperforms the legacy MPI-only algorithm. The quadrant-cache clustering mode was determined to be best suited to the two hybrid algorithms, and is used for the following benchmarks.




Figure 1: Time to solution (x -axis, seconds) for different clustering and memory mode combinations. The results for the 0.5nm benchmark system are shown in the left column, and for the 2.0nm benchmark system in the right column.

The parallel benchmarks for the three algorithms with the 2.0nm on up to 512 Theta nodes are shown in Figure 2. The MPI-only code scales linearly to 256 Xeon Phi nodes, but the memory footprint bottleneck limits utilization of the KNL cores on a given node. This limitation is reduced for the private Fock version of the code, and essentially eliminated for the shared Fock code. On 512 KNL

nodes, the shared Fock code runs about six times faster than the MPI-only implementation.




Figure 2: Multi-node scalability of the three Hartree-Fock algorithms implemented in GAMESS for the 2.0nm benchmark system.

The parallel scaling of the shared Fock code for the 5nm dataset is shown for up to 3,000 Theta KNL nodes in Figure 3. For each KNL node, 4 MPI ranks with 64 OpenMP threads each are utilized. The resulting memory footprint approaches the 208GB memory capacity available per node.




Figure 3: Scalability of the shared Fock MPI/OpenMP algorithm on Theta for the 5.0nm benchmark system.

4 Conclusions

Two hybrid MPI/OpenMP HF codes were implemented in GAMESS. Targeting Theta, the hybrid algorithms significantly reduce the memory footprint by up to 200x, and time-to-solution by up to 6x on KNL and Intel Xeon processors compared to the legacy MPI-only code. Both hybrid codes exhibit better scalability with respect to core counts for both single and multi-node KNL systems, which was demonstrated on as many as 3,000 Theta KNL nodes. These performance improvements extend the scope of chemical problems which can be treated with *ab initio* quantum chemistry methods in GAMESS. Furthermore, the insights gained by this work are applicable to a broader range of codes which solve non-linear partial differential equations using matrix representations.



Argonne Leadership Computing Facility

Argonne National Laboratory
9700 South Cass Avenue, Bldg. #240
Argonne, IL 60439

www.anl.gov



Argonne National Laboratory is a U.S. Department of Energy
laboratory managed by UChicago Argonne, LLC